the globus alliance
www.globus.org

National e-Science Centre

Distributed Monitoring and Information Services for the Grid

Jennifer M. Schopf

Argonne National Laboratory

NeSC

Feb 20, 2006

Argonne National Laboratory, University of Chicago

epcc

KTH Vetenskap och Konst

NCSA

Univa

USC Viterbi School of Engineering

What is a Grid

- Resource sharing
  - Computers, storage, sensors, networks, …
  - Sharing always conditional: issues of trust, policy, negotiation, payment, …
- Coordinated problem solving
  - Beyond client-server: distributed data analysis, computation, collaboration, …
- Dynamic, multi-institutional virtual orgs
  - Community overlays on classic org structures
  - Large or small, static or dynamic

Why is this hard/different?

- Lack of central control
  - Where things run
  - When they run
- Shared resources
  - Contention, variability
- Communication
  - Different sites implies different sys admins, users, institutional goals, and often “strong personalities”

So why do it?

- Computations that need to be done with a time limit
- Data that can't fit on one site
- Data owned by multiple sites

- Applications that need to be run bigger, faster, more

What Is Grid Monitoring?

- Sharing of community data between sites using a standard interface for querying and notification

- A way to discover what services and resources are available to use
- A way to understand the status/attributes of those services
- A system to warn you when things fail

Monitoring Use cases

- PPGD/GriPhyN/iVDGL monitoring group (2002-2004) found roughly 4 categories
  - Health of system (NW, servers, cpus, etc)
  - Resource selection
  - System upgrade evaluation (have systems reached capacity)
  - Application-specific progress tracking
- First three types need roughly the same information
- Fourth is user-specific and application specific – no general solution yet

http://www.mcs.anl.gov/~jms/pg-monitoring

Health of the System "Is the Grid up?"

- Brief Description
  - User of a grid replication service finds actions are much slower than normal
  - Not sure if problem is with network, disk, CPU end points, or something inbetween
  - Need archive data for historical, current streaming for comparison
- Performance events/sensors required
  - Host monitoring - CPU,memory, disk
  - Network path monitoring - bw, lat., traceroute
  - GridFTP monitoring
  - TCP stack monitoring (web 100)
  - Possibly switch/router monitoring
  - May want different data for user vs sys admins

Resource Selection

- Brief Description
  - User/Broker wants to decide where to run a job
  - Sites advertise cluster information for grid-level scheduling decisions
  - Also need data about storage locations and access speeds
  - Information must be summarized for advertising to Grid, scalability is key issue
- Performance events/sensors required
  - Static: number of compute nodes, cpu type and speed, OS, installed sw, available storage systems
  - Dynamic:Queue lengths, large file transfer times

What should monitoring systems look like?

- All sensors must be non-intrusive
- All data is small, and must be “as timely as possible”
- All data must be kept for a long time (years), and must be accessible in many ways
- No one really knows how many sensors will be accessed at one time (or reporting to a higher level service), or how often they will be accessed
- Security isn’t of concern – YET – except for job data

Monitoring Systems (2)

- Line between monitoring system and higher level services isn't always clear
  - Archiving
  - Summary statistics
  - Predictions
  - Error detection
  - Alarms/notification

OUTLINE

- Grid Monitoring and Use Cases
- MDS4
  - Index Service
  - Trigger Service
  - Information Providers
- Deployments
  - Metascheduling data for TeraGrid
  - Service failure warning for ESG
- Performance Numbers

What is MDS4?

- Grid-level monitoring system used most often for resource selection
  - Aid user/agent to identify host(s) on which to run an application
- Uses standard interfaces to provide publishing of data, discovery, and data access, including subscription/notification
  - WS-ResourceProperties, WS-BaseNotification, WS-ServiceGroup
- Part of the Globus Toolkit v4
- Functions as an hourglass to provide a common interface to lower-level monitoring tools

Information Users : Schedulers, Portals, Warning Systems, etc.

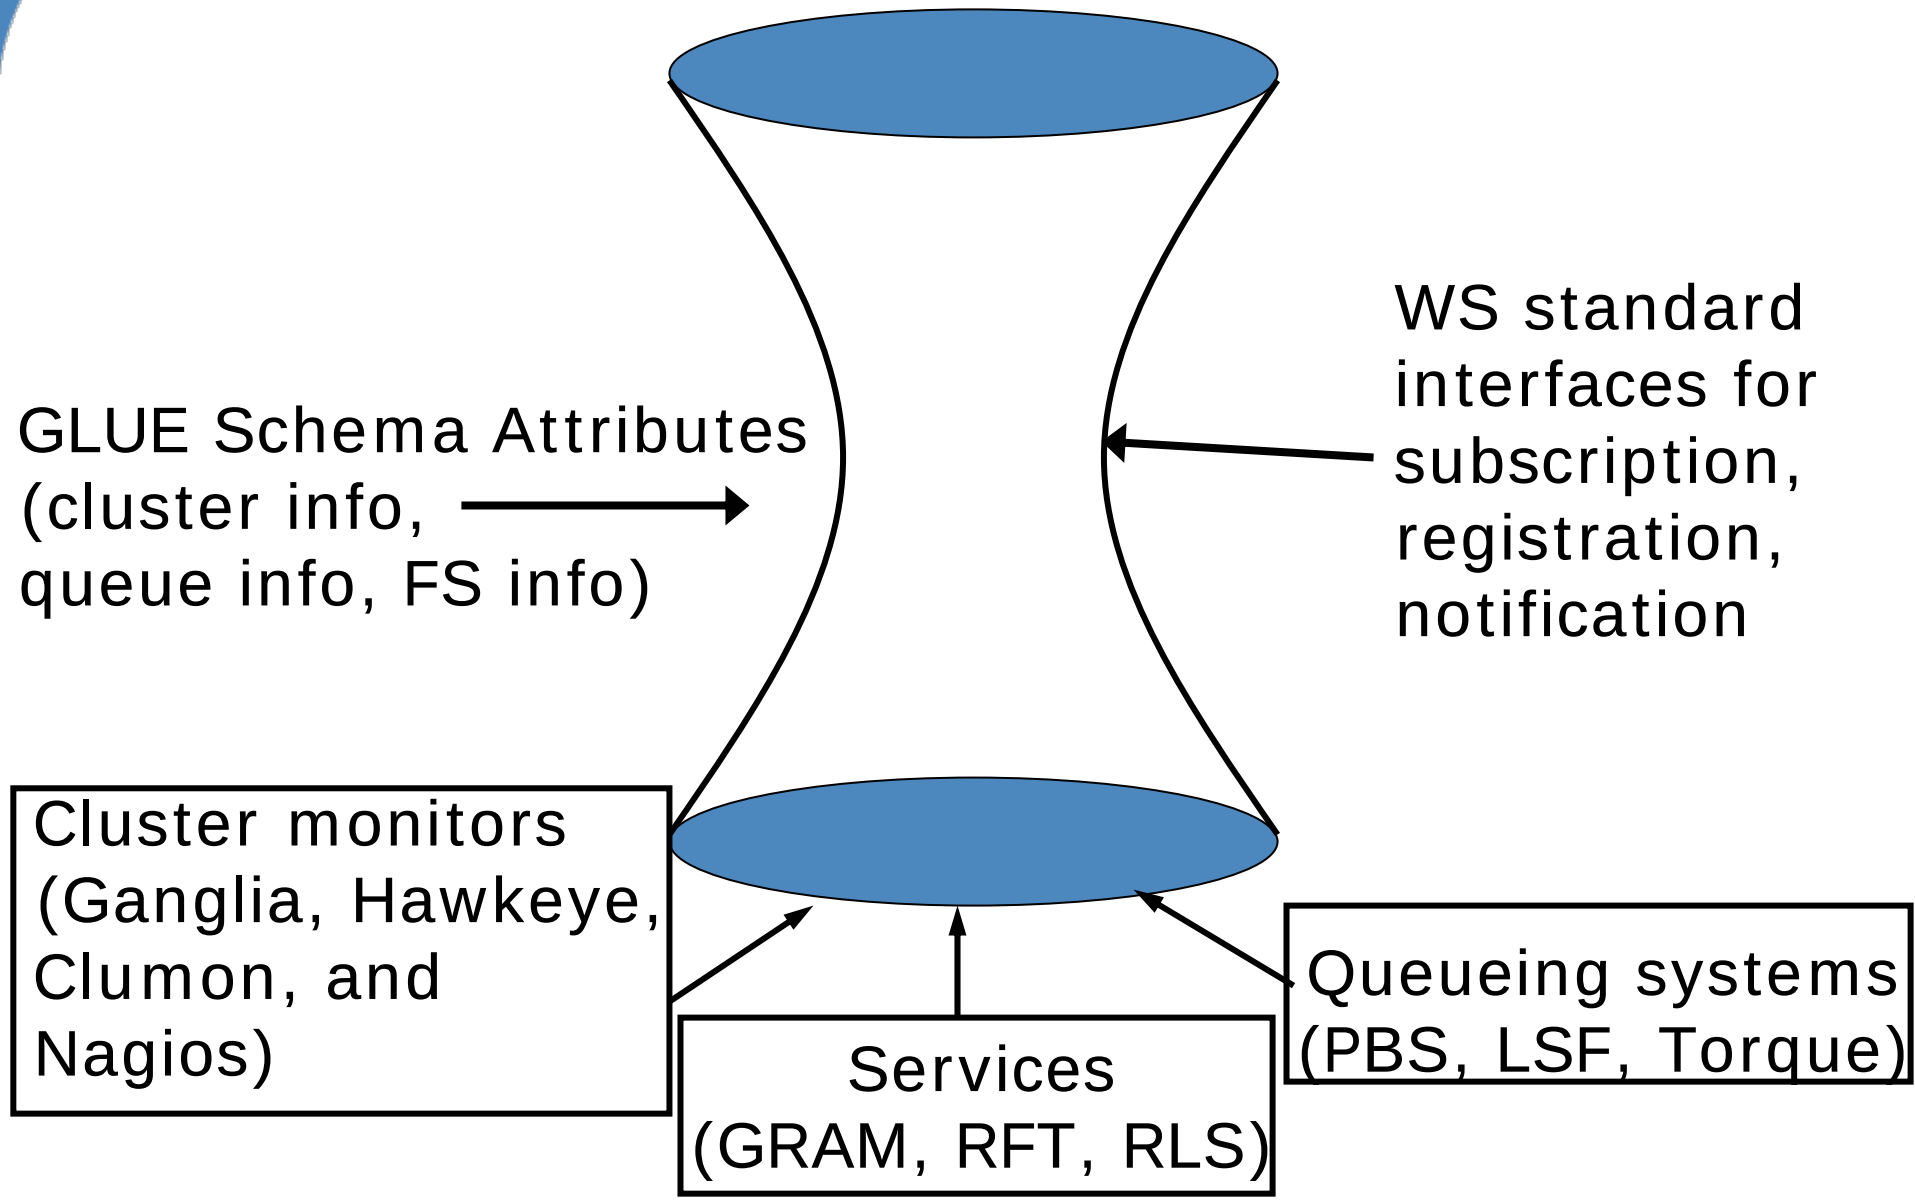

MDS4 Uses Web Service Standards

- WS-ResourceProperties
  - Defines a mechanism by which Web Services can describe and publish resource properties, or sets of information about a resource
  - Resource property types defined in service's WSDL
  - Resource properties can be retrieved using WS-ResourceProperties query operations
- WS-BaseNotification
  - Defines a subscription/notification interface for accessing resource property information
- WS-ServiceGroup
  - Defines a mechanism for grouping related resources and/or services together as service groups

MDS4 Components

- Higher level services
  - Index Service – a way to aggregate data
  - Trigger Service – a way to be notified of changes
  - Both built on common aggregator framework
- Information providers
  - Monitoring is a part of every WSRF service
  - Non-WS services can also be used
- Clients
  - WebMDS

- All of the tool are schema-agnostic, but interoperability needs a well-understood common language

MDS4 Index Service

- Index Service is both registry and cache
- Subscribes to information providers
- Publishes (as resource properties)
  - Datatype and data provider info, like a registry
  - Last value of data, like a cache
- In memory default approach, DB backing store currently being developed to allow for very large indexes
- Soft-state registration
- Can be set up for a site or set of sites, a specific set of project data, or for user-specific data only
- Can be a multi-rooted hierarchy

Index Service Facts 1

- No single global Index provides information about every resource on the Grid
  - No person in the world is part of every VO!
  - Hierarchies or special purpose index's are common
  - Each virtual organization will have different policies on who can access its resources
- The presence of a resource in an Index makes no guarantee about the availability of the resource for users of that Index
  - Ultimate decision about whether to use the resources is left to direct negotiation between user and rsc
  - MDS does not need to keep track of policy information (something that is hard to do concisely)
  - Rscs do not need to reveal their policies publicly

Index Service Facts 2

- MDS has a soft consistency model
  - Published information is recent, but not guaranteed to be the absolute latest
  - Load caused by information updates is reduced at the expense of having slightly older information
  - Free disk space on a system 5 minutes ago rather than 2 seconds ago.
- Each registration into an Index Service is subject to soft-state lifetime management
  - All registrations has expiry times and must be periodically renewed
  - Index is self-cleaning, since outdated entries disappearing automatically

MDS4 Trigger Service

- Subscribe to a set of resource properties
- Evaluate that data against a set of pre-configured conditions (triggers)
- When a condition matches, email is sent to pre-defined address

- Similar functionality in Hawkeye

Aggregator Framework

- General framework for building services that collect and aggregate data
  - Index and Trigger service both use this

1) Common interface implemention
   - Java class that implements an interface to collect XML-formatted data from information providers
   - Implements WS-RP and WS-N for query and subscription
2) Common configuration mechanism
   - Maintain information about which information providers to use and their associated parameters
   - Specify what data to get, and from where
3) Services are self-cleaning
   - Each registration has a lifetime
   - If a registration expires without being refreshed, it and its associated data are removed from the server

Information Providers

- Data sources for the higher level services (eg. Index, Trigger)
- WSRF-compliant service
  - WS-ResourceProperty for Query source
  - WS-Notification mechanism for Subscription source
- Other services/data sources
  - Executable program that obtains data via some domain-specific mechanism for Execution source.

Information Providers: Cluster and Queue Data

- Interfaces to Hawkeye, Ganglia, CluMon
  - Not WS so these are Execution Sources
  - Basic host data (name, ID), processor information, memory size, OS name and version, file system data, processor load data
  - Some condor/cluster specific data
- Interfaces to PBS, Torque LSF queue system
  - Queue information, number of CPUs available and free, job count information, some memory statistics and host info for head node of cluster

Information Providers: GT4 Services

- Every WS built using GT4 core
  - ServiceMetaDataInfo element includes start time, version, and service type name
- Reliable File Transfer Service (RFT)
  - Service status data, number of active transfers, transfer status, information about the resource running the service
- Community Authorization Service (CAS)
  - Identifies the VO served by the service instance
- Replica Location Service (RLS)
  - Note: not a WS
  - Location of replicas on physical storage systems (based on user registrations) for later queries

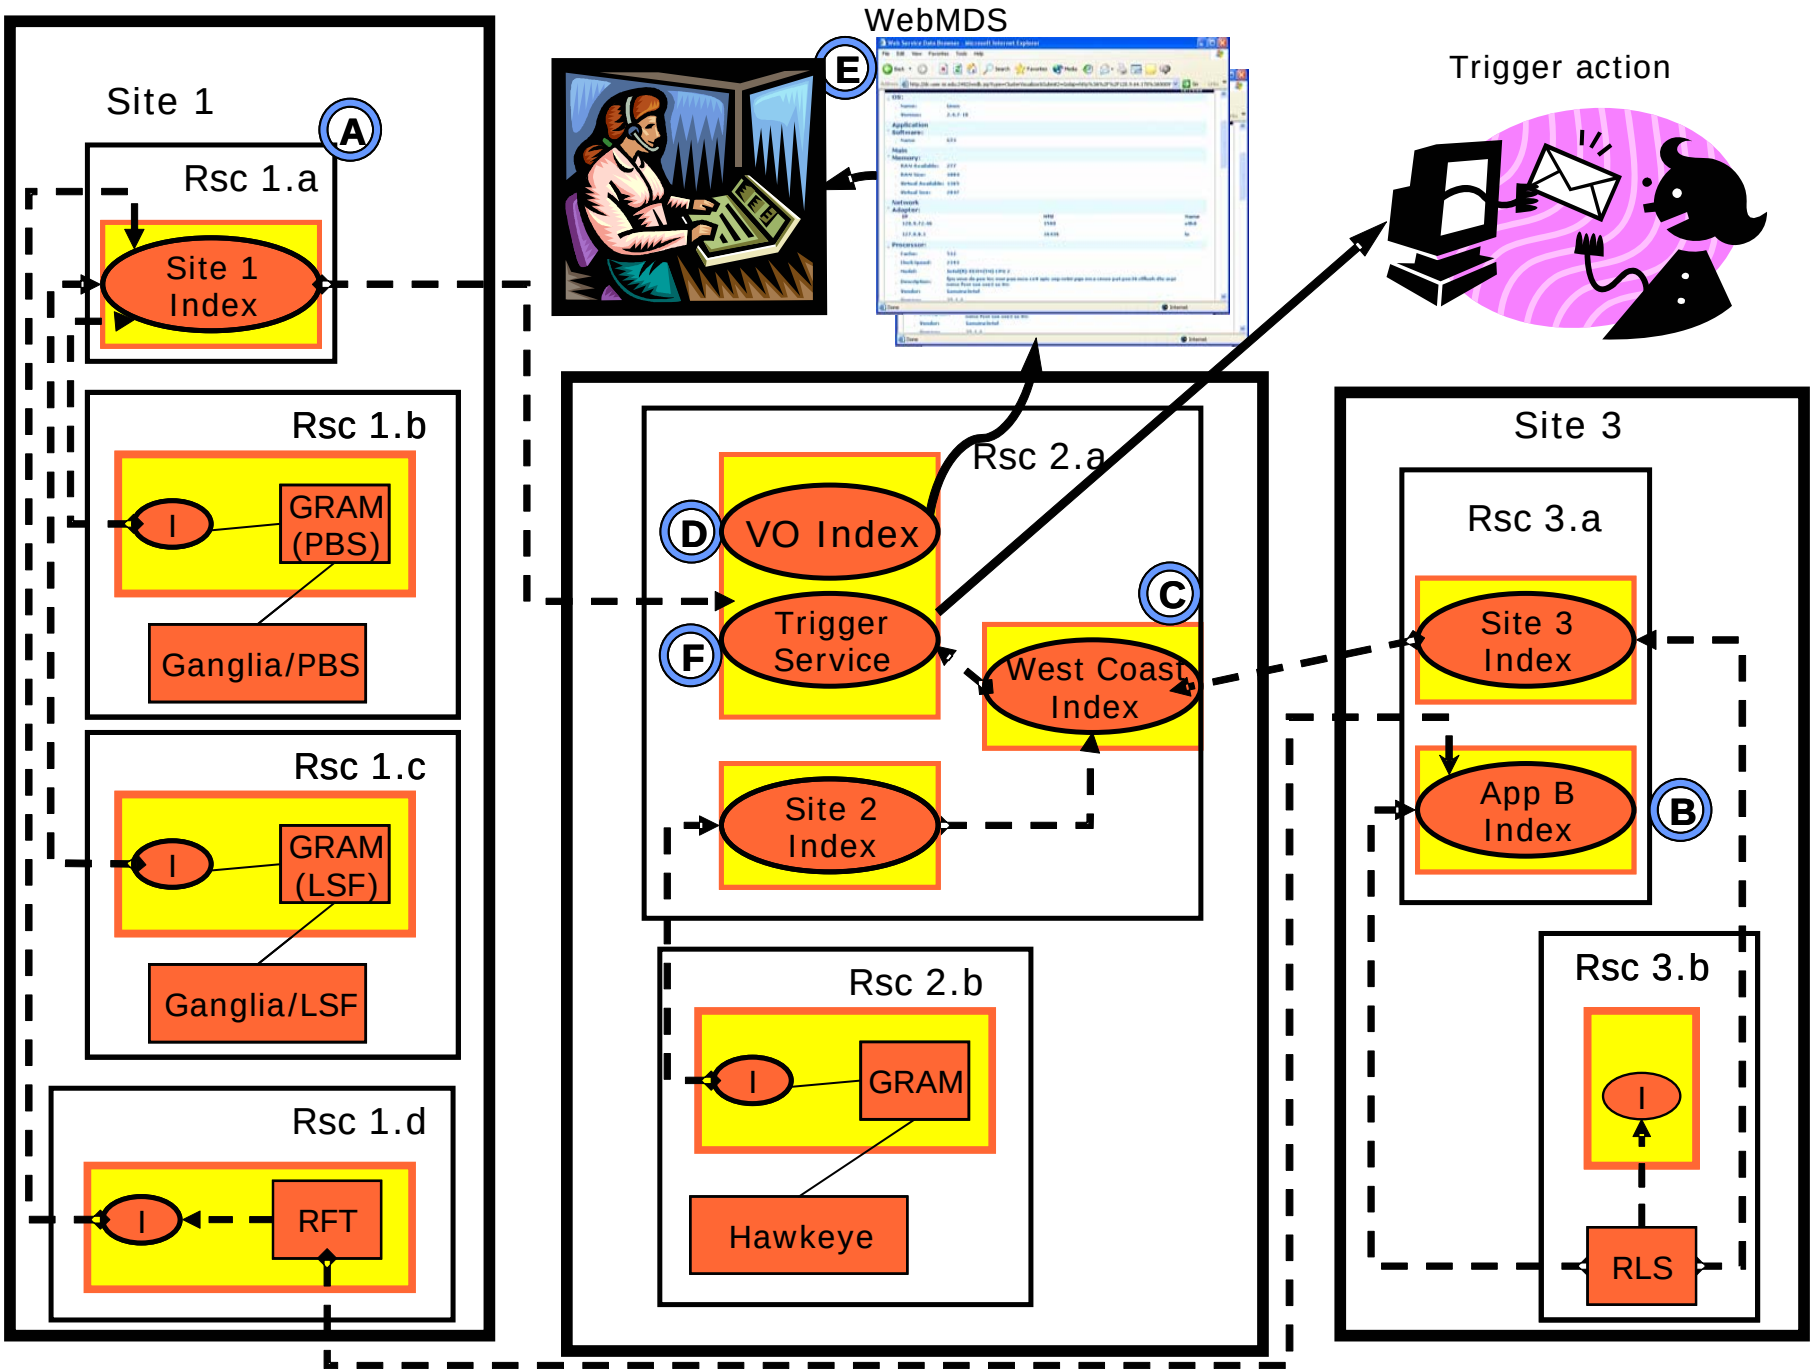
WebMDS
Trigger action
Site 1
A
Rsc 1.a
Site 1 Index
Rsc 1.b
I
GRAM (PBS)
Ganglia/PBS
Rsc 1.c
I
GRAM (LSF)
Ganglia/LSF
Rsc 1.d
I
RFT
E
Rsc 2.a
D
VO Index
F
Trigger Service
C
West Coast Index
Site 2 Index
Rsc 2.b
I
GRAM
Hawkeye
Site 3
Rsc 3.a
Site 3 Index
App B Index
B
Rsc 3.b
I
RLS

WebMDS User Interface

- Web-based interface to WSRF resource property information
- User-friendly front-end to the Index Service
- Uses standard resource property requests to query resource property data
- XSLT transforms to format and display them
- Customized pages are simply done by using HTML form options and creating your own XSLT transforms
- Sample page:
  - http://mds.globus.org:8080/webmds/webmds?info=indexinfo&xsl=servicegroupxsl

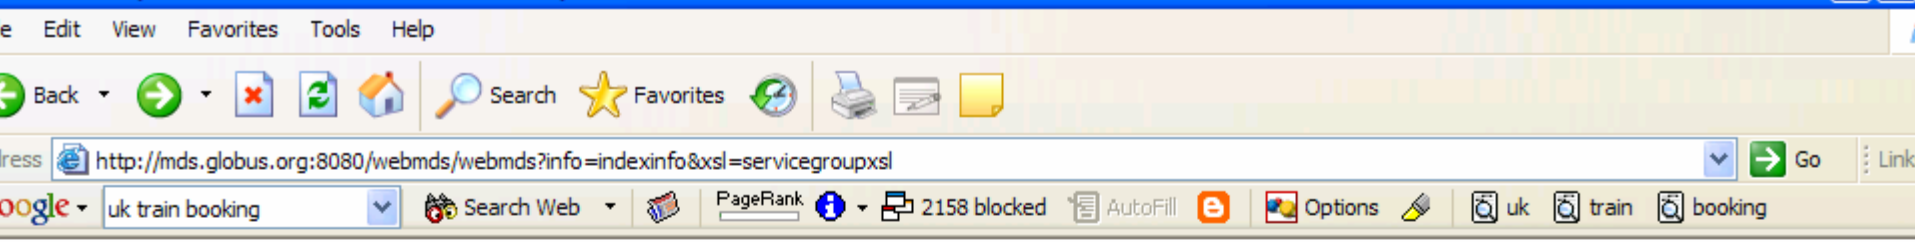

erviceGroup Overview

s page provides a brief overview of Web Services and/or WS-Resources that are members of a WS-ServiceGroup.

s WS-ServiceGroup has 4 direct entries, 33 in whole hierarchy.

esource Type	ID	Information	
Unknown	128.9.72.106	Aggregator entry with no content from `https://128.9.72.106:8443/wsrf/services/ReliableFileTransferFactoryService`	detail
GRAM	128.9.72.106	0 queues, submitting to 0 cluster(s) of 0 host(s).	detail
erviceGroup	128.9.72.140	This WS-ServiceGroup has 11 direct entries, 29 including descendants.	detail
erviceGroup	128.9.72.178	This WS-ServiceGroup has 4 direct entries, 4 including descendants.	detail
RFT	128.9.72.178	0 active transfer resources, transferring 0 files. 40.55 GB transferred in 173769 files since start of database.	detail
GRAM	128.9.72.178	0 queues, submitting to 1 cluster(s) of 10 host(s).	detail
GRAM	128.9.72.178	1 queues, submitting to 1 cluster(s) of 10 host(s).	detail
GRAM	128.9.72.178	2 queues, submitting to 1 cluster(s) of 10 host(s).	detail
erviceGroup	128.9.72.106	This WS-ServiceGroup has 3 direct entries, 3 including descendants.	detail
GRAM	128.9.72.106	0 queues, submitting to 0 cluster(s) of 0 host(s).	detail
GRAM	128.9.72.106	1 queues, submitting to 0 cluster(s) of 0 host(s).	detail
RFT	128.9.72.106	0 active transfer resources, transferring 0 files. 8.28 GB transferred in 8595 files since start of database.	detail
erviceGroup	128.9.64.179	This WS-ServiceGroup has 4 direct entries, 4 including descendants.	detail
GRAM	128.9.64.179	1 queues, submitting to 1 cluster(s) of 15 host(s).	detail
GRAM	128.9.64.179	5 queues, submitting to 1 cluster(s) of 15 host(s).	detail
RFT	128.9.64.179	0 active transfer resources, transferring 0 files. 63.16 GB transferred in 108704 files since start of database.	detail
GRAM	128.9.64.179	0 queues, submitting to 1 cluster(s) of 15 host(s).	detail
erviceGroup	128.9.128.168	This WS-ServiceGroup has 3 direct entries, 3 including descendants.	detail
GRAM	128.9.128.168	0 queues, submitting to 0 cluster(s) of 0 host(s).	detail
RFT	128.9.128.168	0 active transfer resources, transferring 0 files. 10.52 GB transferred in 23489 files since start of database.	detail

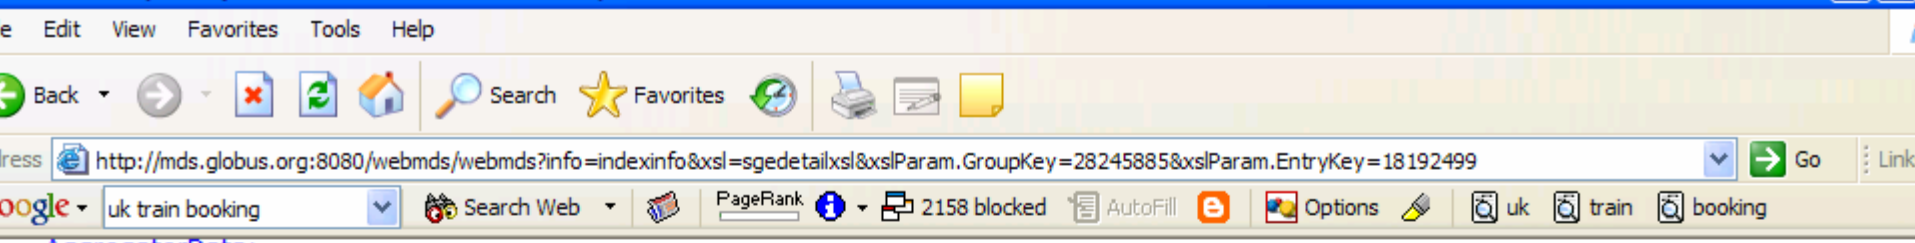

- AggregatorData:
  - GLUECE:
    - Cluster:
      - Name: Viz
      - UniqueID: Viz
      - SubCluster:
        - Name: main
        - UniqueID: main
        - Host:
          - Name: viz-2-int
          - UniqueID: viz-2-int
          - Processor:
            - CacheL1: 0
            - CacheL1D: 0
            - CacheL1I: 0
            - CacheL2: 0
            - ClockSpeed: 2400
            - InstructionSet: x86
          - MainMemory:
            - RAMAvailable: 1390
            - RAMSize: 2026
            - VirtualAvailable: 2408
            - VirtualSize: 3063
          - OperatingSystem:
            - Name: Linux
            - Release: 2.6.11
          - Architecture:
            - SMPSize: 2
          - FileSystem:
            - AvailableSpace: 371293
            - Name: entire-system
            - ReadOnly: false
            - Root: /
            - Size: 404708
          - NetworkAdapter:
            - IPAddress: 192.168.0.171
            - InboundIP: true
            - MTU: 1500
            - Name: viz-2-int
            - OutboundIP: true

- Last5Min: 0

- ComputingElement:
  - Name: normal
  - UniqueID: normal
  - Info:
    - TotalCPUs: 8
  - State:
    - EstimatedResponseTime: 0
    - FreeCPUs: 8
    - RunningJobs: 0
    - Status: enabled
    - TotalJobs: 0
    - WaitingJobs: 0
    - WorstResponseTime: 0
  - Policy:
    - MaxCPUTime: 0
    - MaxRunningJobs: 0
    - MaxTotalJobs: 0
    - MaxWallClockTime: 0
    - Priority: 0
- ComputingElement:
  - Name: low
  - UniqueID: low
  - Info:
    - TotalCPUs: 8
  - State:
    - EstimatedResponseTime: 0
    - FreeCPUs: 8
    - RunningJobs: 0
    - Status: enabled
    - TotalJobs: 0
    - WaitingJobs: 0
    - WorstResponseTime: 0
  - Policy:
    - MaxCPUTime: 0
    - MaxRunningJobs: 0
    - MaxTotalJobs: 0
    - MaxWallClockTime: 0
    - Priority: 0

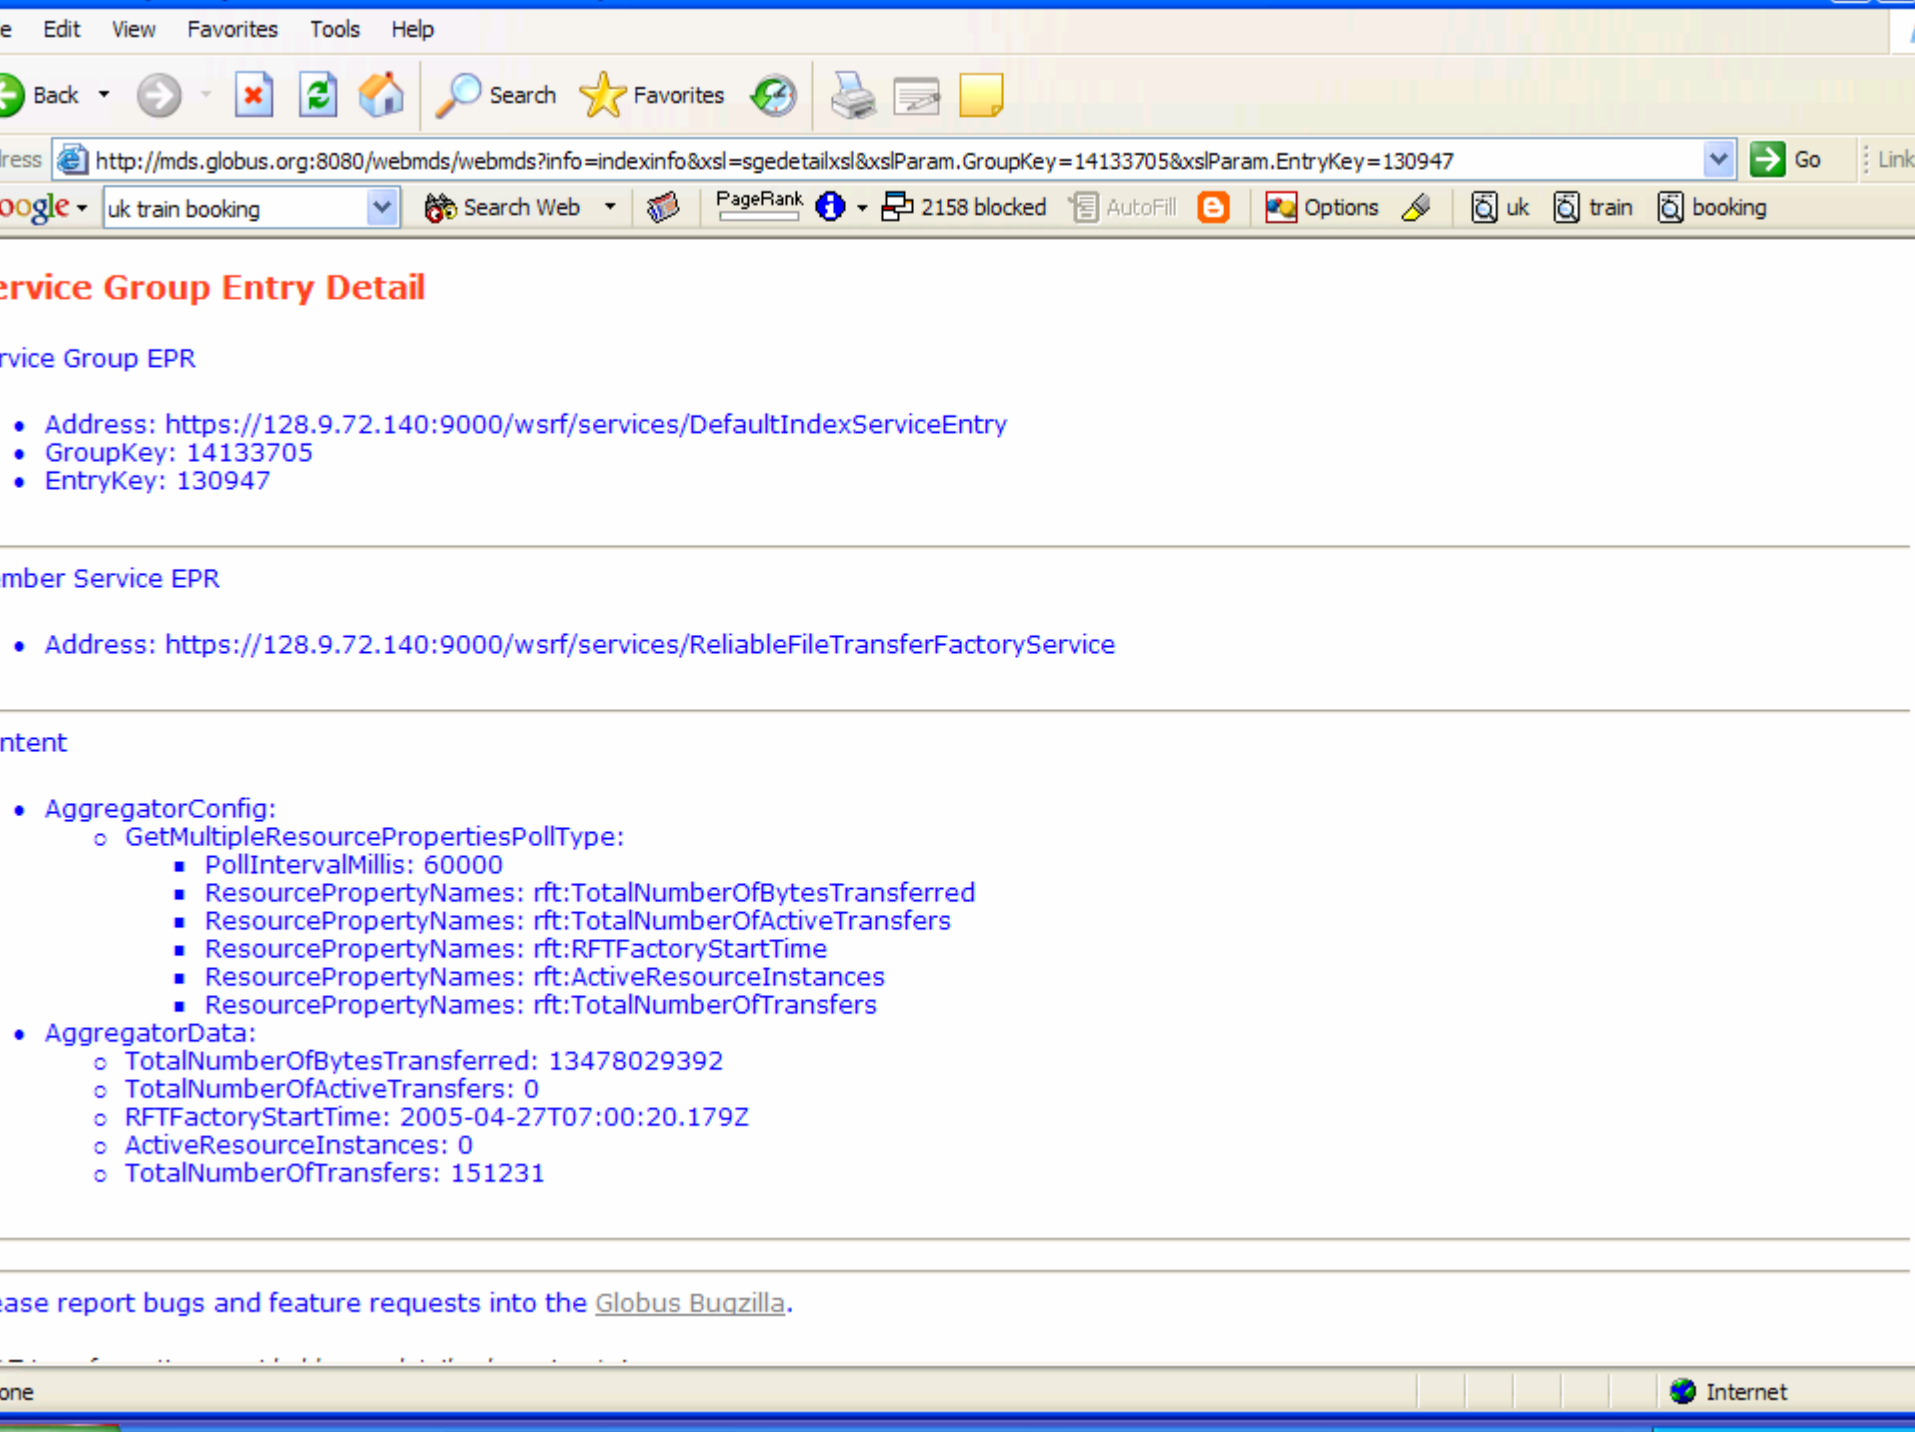

ervice Group Entry Detail

rvice Group EPR

- Address: https://128.9.72.140:9000/wsrf/services/DefaultIndexServiceEntry
- GroupKey: 14133705
- EntryKey: 130947

---

mber Service EPR

- Address: https://128.9.72.140:9000/wsrf/services/ReliableFileTransferFactoryService

---

ntent

- AggregatorConfig:
  - GetMultipleResourcePropertiesPollType:
    - PollIntervalMillis: 60000
    - ResourcePropertyNames: rft:TotalNumberOfBytesTransferred
    - ResourcePropertyNames: rft:TotalNumberOfActiveTransfers
    - ResourcePropertyNames: rft:RFTFactoryStartTime
    - ResourcePropertyNames: rft:ActiveResourceInstances
    - ResourcePropertyNames: rft:TotalNumberOfTransfers
- AggregatorData:
  - TotalNumberOfBytesTransferred: 13478029392
  - TotalNumberOfActiveTransfers: 0
  - RFTFactoryStartTime: 2005-04-27T07:00:20.179Z
  - ActiveResourceInstances: 0
  - TotalNumberOfTransfers: 151231

---

ease report bugs and feature requests into the Globus Bugzilla.

Any questions before I walk through two current deployments?

- Grid Monitoring and Use Cases
- MDS4
  - Index Service
  - Trigger Service
  - Information Providers
- Deployments
  - Metascheduling Data for TeraGrid
  - Service Failure warning for ESG
- Performance Numbers

Working with TeraGrid

- Large US project across 9 different sites
  - Different hardware, queuing systems and lower level monitoring packages
- Starting to explore MetaScheduling approaches
  - GRMS (Poznan)
  - W. Smith (TACC)
  - K. Yashimoto (SDSC)
  - User Portal
- Need a common source of data with a standard interface for basic scheduling info

Cluster Data

- Provide data at the subcluster level
  - Sys admin defines a subcluster, we query one node of it to dynamically retrieve relevant data
- Can also list per-host details
- Interfaces to Ganglia, Hawkeye, CluMon, and Nagios available now
  - Other cluster monitoring systems can write into a .html file that we then scrape

Cluster Info

- UniqueID
- Benchmark/Clock speed
- Processor
- MainMemory
- OperatingSystem
- Architecture
- Number of nodes in a cluster/subcluster
- TG specific Node properties
- StorageDevice
  - Disk names, mount point, space available

Data to collect: Queue info

- Interface to PBS (Pro, Open, Torque), LSF

- LRMSType
- LRMSVersion
- DefaultGRAMVersion and port and host
- TotalCPUs
- Status (up/down)
- TotalJobs (in the queue)
- RunningJobs
- WaitingJobs
- FreeCPUs
- MaxWallClockTime
- MaxCPUTime
- MaxTotalJobs
- MaxRunningJobs

How will the data be accessed?

- Java and command line APIs to a common TG-wide Index server
  - Alternatively each site can be queried directly
- One common web page for TG
  - http://snipurl.com/j24r
- Query page is next!

quake	quake	4.0.1	tg-login1.ncsa.teragrid.org	PBS	2.0.0p0-snap.1129668328	891	533	enabled 0	0	0	2880	-1	-1	-1
gpfs-wan	gpfs-wan	4.0.1	tg-login1.ncsa.teragrid.org	PBS	2.0.0p0-snap.1129668328	891	533	enabled 3	0	3	1440	-1	-1	-1
lq	lq	4.0.1	tg-login2.sdsc.teragrid.org	PBS	1.2.0p6	262	210	enabled 2	1	1	12000	-1	-1	-1
dque	dque	4.0.1	tg-login2.sdsc.teragrid.org	PBS	1.2.0p6	262	210	enabled 54	11	43	1080	-1	-1	-1

Cluster / Subcluster Overview

Type	Name	UniqueID	Processor		Total Memory	Operating System	SMP Size	TeraGrid Extensions
			Type	Clock Speed				Total Nodes
Cluster	NCSA-TeraGrid	NCSA-TG						891
SubCluster	NCSA-TG-IA64CPU13-FASTIO-HIMEM	IA64CPU13-FASTIO-HIMEM	IA-64	1296	4056	Linux2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	128
SubCluster	NCSA-TG-IA64CPU13-FASTIO-LOMEM	IA64CPU13-FASTIO-LOMEM	IA-64	1296	4101	Linux2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	128
SubCluster	NCSA-TG-IA64CPU15-FASTCPU-GPFSWAN	IA64CPU15-FASTCPU-GPFSWAN	IA-64	1496	4106	Linux2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	16
SubCluster	NCSA-TG-IA64CPU15-FASTCPU	IA64CPU15-FASTCPU	IA-64	1496	4106	Linux2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	615
SubCluster	NCSA-TG-IA64CPU13-FASTIO-HIMEM-SPARE	IA64CPU13-FASTIO-HIMEM-SPARE	IA-64	1296	4056	Linux2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	1
SubCluster	NCSA-TG-IA64CPU13-FASTIO-LOMEM-SPARE	IA64CPU13-FASTIO-LOMEM-SPARE	IA-64	1296	4101	Linux2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	1
SubCluster	NCSA-TG-IA64CPU15-PHASE2-FASTCPU-SPARE2	IA64CPU15-PHASE2-FASTCPU-SPARE2	IA-64	1496	4106	Linux2.4.21.SuSE_292.til#1 SMP Fri Jun 3 07	2	2

Hosts in Subcluster NCSA-TG-IA64CPU13-FASTIO-HIMEM

Name	UniqueId	Storage Device				TeraGrid Extensions
		Name	Size	Available Space	Transfer Rate	Node Properties
tg-c001	tg-c001	entire-system	352829	177270	0	all,ia64-compute,compute,ia64-cpu13,fastio,himem,rack40,clos12,stage
tg-c002	tg-c002	entire-system	352816	177248	0	all,ia64-compute,compute,ia64-cpu13,fastio,himem,rack40,clos12
tg-c003	tg-c003	entire-system	352816	177248	0	all,ia64-compute,compute,ia64-cpu13,fastio,himem,rack40,clos12

Status

- Currently have a demo system up
  - Queuing data from SDSC and NCSA
  - Cluster data using CluMon interface at NCSA
  - Basic WebMDS interface
- Being deployed more widely for TeraGrid this week
- General patch for 4.0.1 deployments should be available soon – let me know if you're interested!

ESG use of MDS4 Trigger Service

- Need a way to notify system administrators and users what the status of their services are
- In particular, interested in
  - Replica Locatoin Service (RLS)
  - Storage Resource Manager service (SRM)
  - OpenDAP
  - Web Server (HTTP)
  - GridFTP fileservers

Trigger Service and ESG Cont.

- The Trigger service periodically checks to see if services are up and running
- If a service is gone down or is unavailable for any reason, an action script is executed
  - Sends email to administrators
  - Update portal status page
- Been in use for over a year (used GT3 version previously)

ESG Current Status

Updated: Fri Nov 4 12:00:01 2005 MDT

	LANL	LBNL	NCAR	ORNL
MSS/HPSS				
SRM				
RLS				
OpenDAPg				
GridFTP server				
HTTP server				

(Explanation of current status)

OUTLINE

- Grid Monitoring and Use Cases
- MDS4
  - Index Service
  - Trigger Service
  - Information Providers
- Deployments
  - Metascheduling Data for TeraGrid
  - Service Failure warning for ESG
- Performance Numbers

MDS4 Stability

Vers.	Index Size	Time up (Days)	Queries Processed	Query Per Sec.	Round-trip Time (ms)
4.0.1	25	66+	81,701,925	14	69
4.0.1	50	66+	49,306,104	8	115
4.0.1	100	33	14,686,638	5	194
4.0.0	1	14	93,890,248	76	13
4.0.0	1	96	623,395,877	74	13

Scalability Experiments

- MDS index
  - Dual 2.4GHz Xeon processors, 3.5 GB RAM
  - Sizes: 1, 10, 25, 50, 100
- Clients
  - 20 nodes also dual 2.6 GHz Xeon, 3.5 GB RAM
  - 1, 2, 3, 4, 5, 6, 7, 8, 16, 32, 64, 128, 256, 384, 512, 640, 768, 800
- Nodes connected via 1Gb/s network
- Each data point is average of 8 minutes
  - Ran for 10 mins but first 2 spent getting clients up and running
  - Error bars are SD over 8 mins
- Experiments by Ioan Raicu, U of Chicago, using DiPerf

MDS4 Throughput

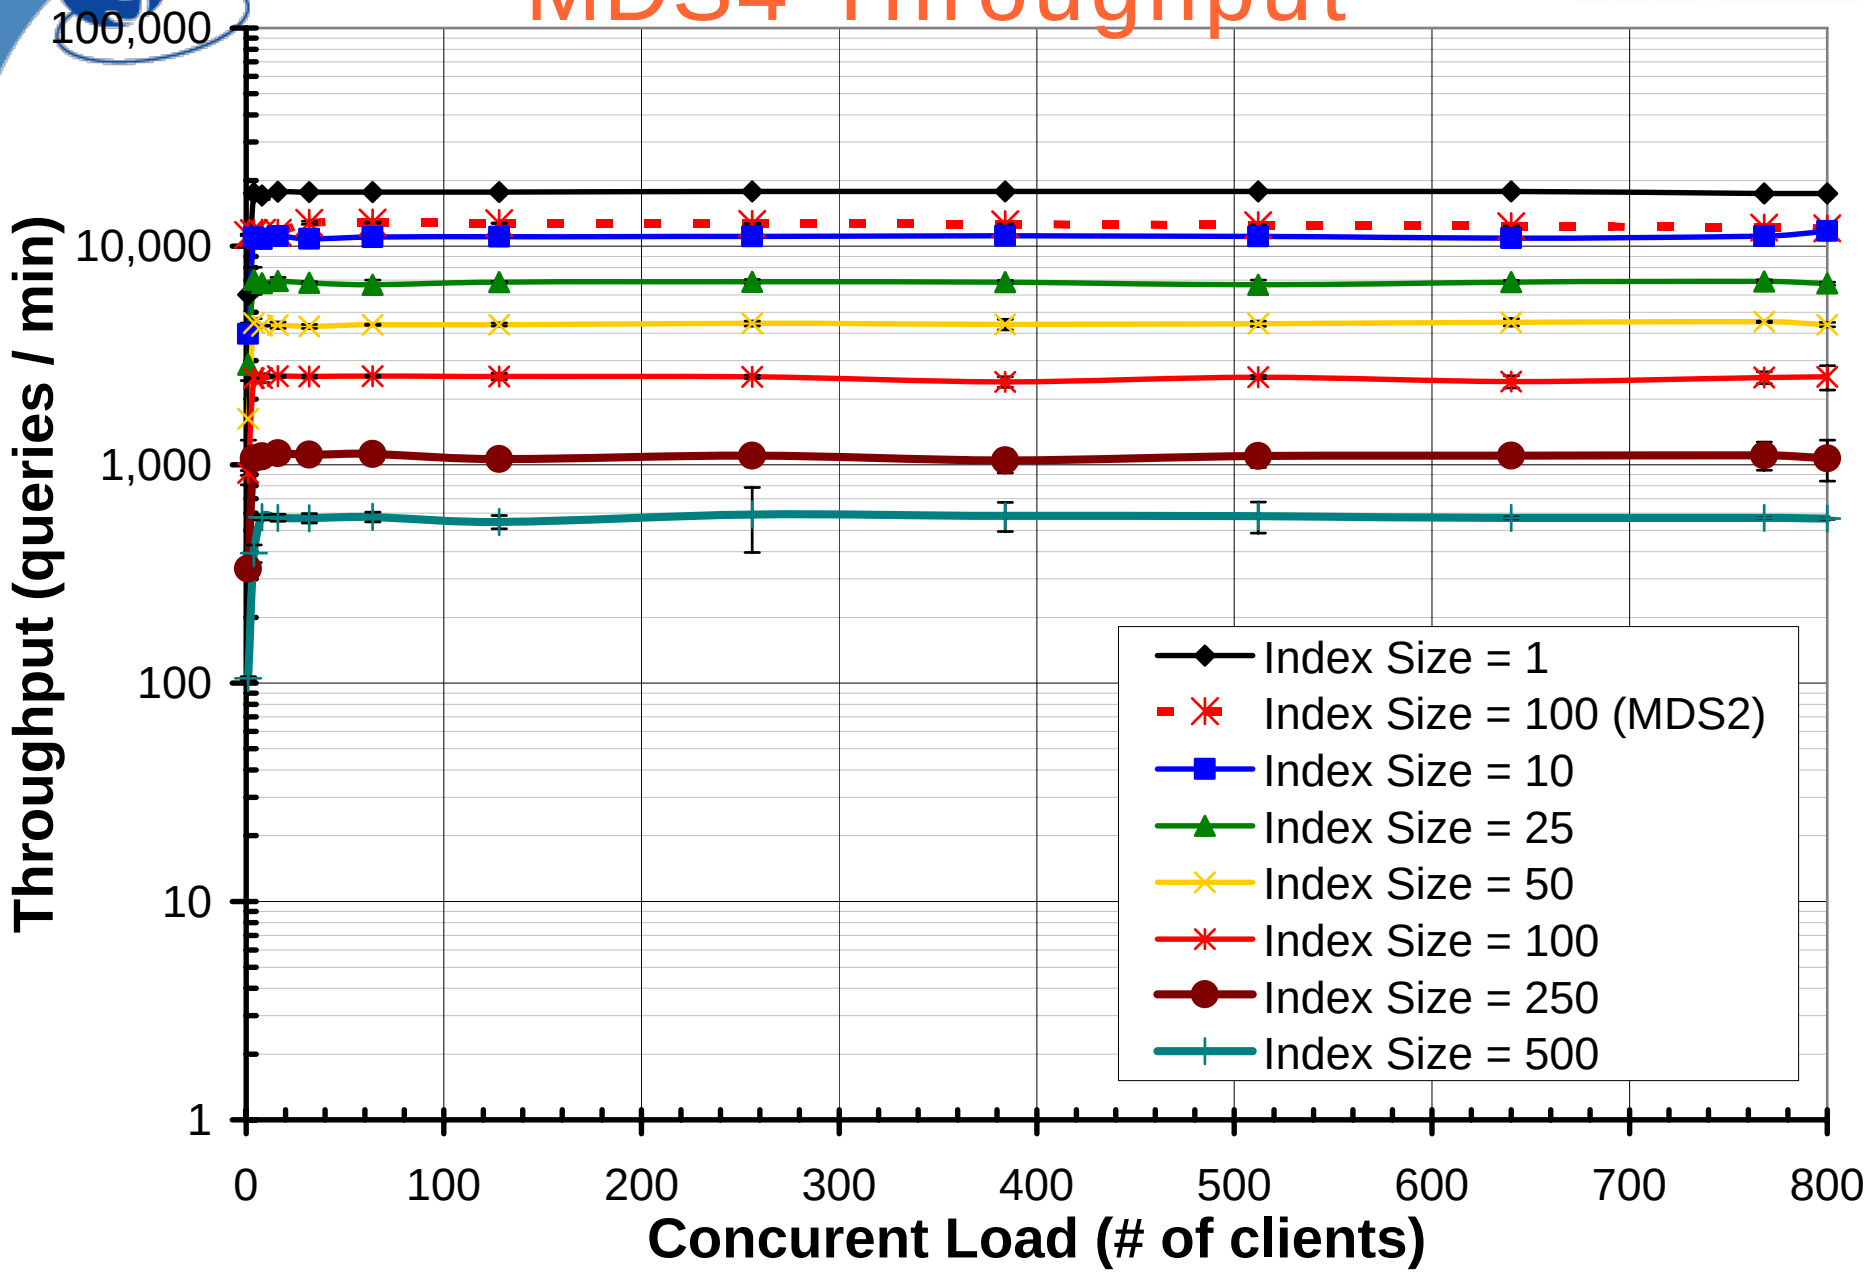

MDS4 Response Time

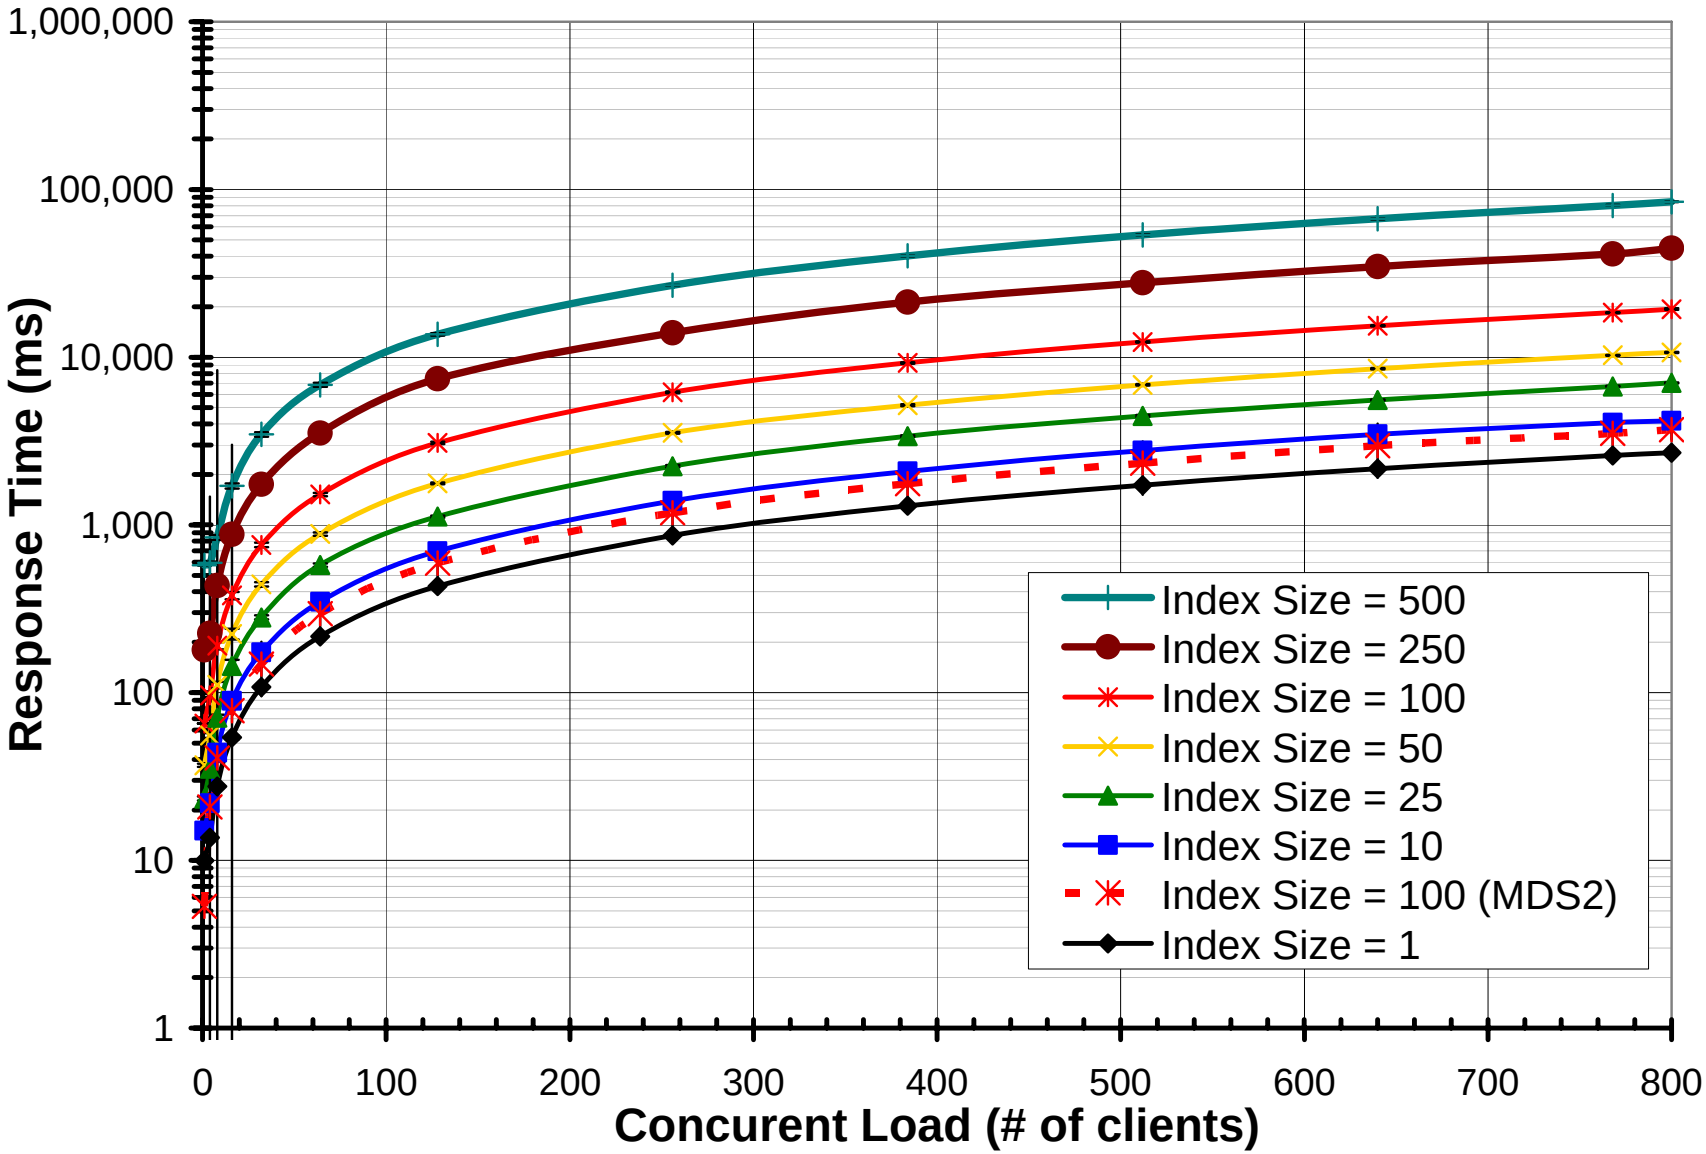

Index Maximum Size

Heap Size (MB)	Approx. Max. Index Entries	Index Size (MB)
64	600	1.0
128	1275	2.2
256	2650	4.5
512	5400	9.1
1024	10800	17.7
1536	16200	26.18

Performance

- Is this enough?
  - We don't know!
  - Currently gathering up usage statistics to find out what people need
- Bottleneck examination
  - In the process of doing in depth performance analysis of what happens during a query
  - MDS code, implementation of WS-N, WS-RP, etc

(These numbers are in an HPDC submission)

Summary

- MDS4 is a WS-based Grid monitoring system that uses current standards for interfaces and mechanisms
- Available as part of the GT4 release
  - Currently in use for resource selection and fault notification
- Initial performance results aren't awful – we need to do more work to determine bottlenecks

Where do we go next?

- Extend MDS4 information providers
  - More data from GT4 WS
    - GRAM, RFT, CAS
  - More data from GT4 non-WS components
    - RLS, GridFTP
  - Interface to other data sources
    - Inca, GRASP
  - Interface to archivers
    - PinGER, NetLogger
- Additional scalability testing and development
- Additional clients

Other Possible Higher Level Services

- Archiving service
  - The next high level service we'll build
  - Looking at Xindice as a possibility
- Site Validation Service (ala Inca)
- Prediction service (ala NWS)
- What else do you think we need?

Contributing to MDS4

- Globus is opening up it's development environment – similar to Apache Jakarta
- MDS4 will be a project in the new scheme
- Contact me for more details
  - jms@mcs.anl.gov
- http://dev.globus.org

Thanks

- MDS4 Team: Mike D'Arcy (ISI), Laura Pearlman (ISI), Neill Miller (UC), Jennifer Schopf (ANL)
- Students: Ioan Raicu, Xuehai Zhang
- This work was supported in part by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under contract W-31-109-Eng-38, and NSF NMI Award SCI-0438372. This work also supported by DOESG SciDAC Grant, iVDGL from NSF, and others.

For More Information

- Jennifer Schopf
  - Jms@mcs.anl.gov
  - http://www.mcs.anl.gov/~jms
- Globus Toolkit MDS4
  - http://www.globus.org/toolkit/mds
- Monitoring and Discovery in a Web Services Framework: Functionality and Performance of the Globus Toolkit's MDS4
  - http://www.mcs.anl.gov/~jms/Pubs/mds4.hpdc06.pdf